

# ASCR

ADVANCED SCIENTIFIC  
COMPUTING RESEARCH

## EXASCALE REQUIREMENTS REVIEW

An Office of Science review sponsored by  
Advanced Scientific Computing Research

September 27-29, 2016  
Rockville, Maryland



U.S. DEPARTMENT OF  
**ENERGY**

# ADVANCED SCIENTIFIC COMPUTING RESEARCH



1

YEAR OF  
PLANNING



3

DAYS OF  
DISCUSSION



77

NEEDS  
IDENTIFIED:  
52 CLASSIFIED  
AS HIGH  
PRIORITY



130

PARTICIPANTS  
FROM NEARLY  
30 INSTITUTIONS

## Exascale Requirements Reviews: Overview

During 2015 and 2016, the U.S. Department of Energy (DOE) conducted Exascale Requirements Reviews for each of its six Office of Science (SC) program offices. The goal of the reviews was to help ensure the ability of DOE's Advanced Scientific Computing Research (ASCR) facilities to support SC mission science in the exascale age. The ASCR review brought together researchers and facilities planners to identify the systems, facilities, processes, software, support, and policies required to create a productive exascale computing ecosystem for ASCR research in applied mathematics, computer science, and networking over the next decade.

DOE program managers are using the review reports to guide strategic planning and investments for the 2020–2025 time frame.

## ASCR Grand Challenges


The widespread use of computing in the American economy, and especially in science, would not be possible without a thoughtful, exploratory research and development (R&D) community advancing foundational components such as algorithms, operating systems, computer languages, software libraries, and tools for data management and analysis. These are the tools and building blocks—the hammers, chisels, bricks, and mortar—of the smartphone, the cloud, and the computing services on which we all rely, but they are especially important for the very large, complex, high-performance computers supported by the ASCR Facilities Division. Engineers and scientists need ever-more-capable computing tools to discover new material properties for manufacturing, make energy generation safer and more efficient, and provide insight into the fundamentals of the universe. The research division of ASCR (ASCR Research) ensures that these tools and building blocks are being developed and honed to meet the needs of modern computational and data-intensive science.

**The information captured in a pre-review survey and in the breakout sessions revealed that the needs of the ASCR Research community differ from those of other DOE SC programs because of ASCR's focus on leading-edge research in computer science, applied mathematics, and networking.**

## Answering ASCR Challenges in the Exascale Age

The high-performance computing (HPC) environment is a computational ecosystem with diverse, complex, and tightly inter-related components. Enhancing that ecosystem benefits all aspects of the ASCR Research program and the other SC program offices. The needs identified by ASCR researchers varied across the spectrum from emerging architectures to early delivery systems to production systems, but a number of common themes emerged. The findings are not presented in order of priority—each of them includes needs identified as “high-priority”; they should be viewed as an interrelated set of needs that combine to form the key findings of the review.

- In contrast to the domain science areas, the ASCR Research focus on computer science, applied mathematics, and next-generation networking means that current and future HPC systems themselves are frequently the subject of the research. Thus, researchers need *frequent, comprehensive access* to these systems, as well as a better understanding of the R&D required to make the systems usable and useful for science applications.
- Close collaboration and effective communication between HPC facilities and ASCR researchers are critical to facilitating researcher access to HPC resources, as well as enabling researchers to *study and test HPC systems in detail*.
- Because HPC facilities are often the venue for ASCR researchers to develop, test, and deploy their tools, these facilities need to provide a computing ecosystem that supports a *development environment for ASCR research efforts*, as well as an environment for other scientists who are HPC users. Software tools produced by ASCR researchers enable physicists, biologists, and other domain scientists to fully exploit the power of machines at the ASCR facilities and thereby increase their scientific output.



Computational scientists face challenges in deciding which data are the most essential for analysis and transforming these data into meaningful visual representations. This image was generated using the information-theoretic streamline placement algorithm (Image courtesy of Han-Wei Shen, Ohio State University; Rob Ross and Tom Peterka, Argonne; and Yi-Jen Chiang, Polytechnic Institute of New York University).

# ADVANCED SCIENTIFIC COMPUTING RESEARCH

## ASCR Research Programs

The mission of the ASCR program is to advance applied mathematics and computer science; deliver the most advanced computational scientific applications in partnership with disciplinary science; advance computing and networking capabilities; and develop future generations of computing hardware and software tools for science, in partnership with the research community, including U.S. industry.

The *Computer Science* program pursues innovative advances in programming languages and environments; system software; computer architectures; performance and productivity tools; networks; and data management, analysis, and visualization; among many other areas. The primary focus is on effective use of very large-scale computers, networks, and data.

The *Applied Mathematics* program supports mathematical and computational research that facilitates the use of the latest HPC systems to advance our understanding of science and technology.

*Scientific Discovery through Advanced Computing (SciDAC)* brings together experts from across application domains and from universities, national laboratories, and industry to ensure that scientists are using state-of-the-art technologies to solve their increasingly complex computational and data science challenges.

For computer science researchers, computing facilities are not just a means to an end—the software required to use them effectively is the subject of the research. Similarly, while other SC researchers rely on high-performing networks to carry their data, networking researchers study the detailed performance of the network itself. Finally, SC scientists rely on efficient, scalable algorithms embedded in their simulation codes or available in software libraries at the facilities, and ASCR applied mathematicians are at the cutting edge of defining new algorithms to run on new machines.

ASCR researchers' output ranges from papers to increase fundamental understanding of hardware systems, software and programming systems, and numerical algorithms to the design and development of fully functioning production software that enables cutting-edge domain science.

The Exascale Requirements Review reports and supporting materials can be found at <http://exascale.org>

DOE's HPC centers are based at Argonne National Laboratory (ALCF), Lawrence Berkeley National Laboratory (NERSC), and Oak Ridge National Laboratory (OLCF), with

networking services provided by ESNet; all of these facilities operate under the direction of ASCR.